



(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
18.09.1996 Bulletin 1996/38

(51) Int Cl.⁶: **G06F 13/28**

(21) Application number: 96301799.1

(22) Date of filing: 15.03.1996

(84) Designated Contracting States:
DE FR GB

(30) Priority: 17.03.1995 US 407439

(71) Applicant: SYMBIOS LOGIC INC.
Fort Collins, Colorado 80525 (US)

(72) Inventors:
• Carmichael, R.
Longmont, CO 80503 (US)


- Ward, J.M.
Fort Collins, CO 80525 (US)
- Winchell, M.A.
Fort Collins, CO 80525 (US)

(74) Representative: Gill, David Alan
W.P. Thompson & Co.,
Celcon House,
289-293 High Holborn
London WC1V 7HU (GB)

(54) **Controlling (n+1) I/O channels with (n) data managers in a homogeneous software programming environment**

(57) The invention provided for a apparatus and method for improving the input/output performance of a computer system under the control of a multi-tasking, multi-threaded operating system. In particular, the invention provides an apparatus and method to chain contiguous DMA scatter gather sub blocks of a physical region descriptor table for channel 0 with contiguous DMA scatter gather sub blocks of physical region descriptor table for channel 1, using a single data manager, while maintaining maximum media bandwidth. DMA block transfers are scheduled based on the availability of data from the buffer memory of input/output devices (127, 129), thus minimizing both media or network idle time as well as minimizing input/output bus idle time. Near maximum aggregate bandwidth of multiple input/output buses and their associated devices is then obtained, and the apparatus and method thus provides significant performance advantages over prior techniques having two input/output channel systems implemented with a single data manager.

FIG. 1



Description

The present invention relates to a method and apparatus for controlling (n+i) I/O channels with (n) data managers in a homogeneous software programming environment.

In particular the invention can relate to computer system input output (I/O) operations, and more particularly to an apparatus and method for improving I/O performance of a multiple I/O channel computer system at minimal cost.

In a particular embodiment, the invention is intended for use in computer systems controlled by a multiple tasking, multiple threaded operating system such as Windows NT, OS/2, System 7.5, Novell or UNIX, but can also provide significant cost advantages in computer systems controlled by single threaded and/or single tasking operating systems such as Windows or DOS.

A typical computer system generally includes one or more central processing units (CPUs), main memory and one or more I/O channels. The functions of the I/O channels are to transfer data between the CPU's main memory and either storage units or network interface devices. Storage units store data and programs which the CPU uses in performing specific programming tasks. Typical storage units include hard disk drives, CDROMs and tape drives. Network interface devices allow the transfer of data to or from other computer systems on the network. Typical network interfaces include ethernet, fiber channel, ATM and FDDI.

Recently, applications run by the CPU have moved towards using the high-performance processing power of microprocessors (such as the Intel Pentium, PowerPC and MIPS R4400) to generate real-time, full motion video. To support these applications, substantially more I/O bandwidth is required. Typically, full motion video windows are limited in size to a fraction of the video monitor size due to limitations in I/O bandwidth and or limitations in the processing power of the CPU or video card. Users operating computer systems in multi-window, graphics intensive environments expect instant response to mouse clicks or commands. Mouse clicks on an ICON representing a large text file or application should result in immediate visual results, whether or not the selected file resides on an I/O channel, or whether or not an I/O device is currently transferring large blocks of full motion video data. This instant response expectation, when coupled with the high-bandwidth transfer requirement, dictates that future I/O controllers allow I/O requests to be preempted and rescheduled at a later time.

More recently, new techniques have been proposed for providing low-cost, high-performance I/O in the personal computer (PC) area. Enhanced IDE disk drives having large buffers, capable of interfacing to IDE channels having over sixteen megabyte per second DMA bandwidth, have been demonstrated and produced, and will likely become mainstream in the PC industry over

the next few years. Brad Hosler has proposed an IDE interface specification for improving system performance in multitasking environments, called "A Programming Interface for Bus Master IDE Controllers", which assumes the capabilities inherent in the DMA capable IDE disk drives mentioned earlier. Hosler's proposal specifies a simple scatter/gather mechanism which allows large or small blocks of I/O transfer data to be scattered to or gathered from main memory. This mechanism cuts down on the number of processor interrupts required and the total number of I/O device interactions with the CPU when transferring large blocks of data, such as that required to support full motion video. Although the scatter gather programming interface specifications were originally intended specifically for controlling hard disk drives on an IDE channel, software drivers can be written to access other types of storage devices on the IDE channel, storage devices on other types of I/O channels or network interface devices using the same homogeneous scatter gather programming interface.

Fig.1 shows a known apparatus used to perform bus master scatter gather DMA operations. An I/O control device 20 contains one or more descriptor table pointer registers 22, a data manager 24, an I/O interface 26 and a CPU interface 28. Each data manager 24 contains (i) a dedicated DMA engine for managing bus master data transfers between the I/O device (not shown) and main memory 30, and (ii) buffer memory required for maintaining maximum data transfer bandwidth between the I/O channel or interface bus 32 and the processor interface 28. The I/O control device 20 (also known as an I/O controller) connects to the CPU 34 and main memory 30 through a bridge 36, via an interface 21. The bridge 36 may simply provide an extension of the processor's bus, or may buffer and extend the processor bus using an entirely different bus structure and protocol such as PCI. These types of bridges are commonly known in the art.

Besides containing CPU program and data storage, main memory 30 contains one or more physical region descriptor tables 37 having entries 38. The main memory 30 allows storage of blocks of DMA data 40 transferred to or from the I/O devices on the I/O channel or interface bus 32. Data regions 40 represent blocks of linear DMA transfers of contiguous double words of data, as specified by a respective entry 38 in the physical region descriptor table 37.

Fig 2 shows a prior method used for performing I/O operations. At step 42, the CPU prepares a physical region descriptor (PRD) table in main memory. Each PRD table entry is 8 bytes long and consists of an address pointer to the starting address and the transfer size of the memory buffer to be transferred. The PRD table may contain up to 8192 PRD table entries, with each entry specifying a single, contiguous block of data to be transferred between the I/O channel device and main memory. Sequential PRD table entries may specify that a

block of data be transferred to or from any available location in main memory.

At step 44, the CPU writes the starting address of the main memory physical region descriptor table to the I/O control device's descriptor table pointer register. The CPU writes the I/O controller's DMA transfer direction bit and clears appropriate I/O controller interrupt bits and error flags resulting from previous transfers.

At step 46, the CPU issues the appropriate DMA command to the I/O device.

At step 48, the CPU writes a start scatter gather command to the I/O controller.

At steps 50, 52 and 54, the I/O controller performs a bus master scatter gather DMA operation by sequentially performing all contiguous DMA block transfers between the I/O device and main memory as specified by the complete set of entries within a PRD table.

At step 56, the I/O device signals that the requested data transfer is complete. The I/O controller provides signals the processor when all of the requested data has travelled through the I/O controller's internal FIFOs and has been written successfully to either the main memory or the I/O device.

The I/O controller is usable for the next bus master DMA operation.

Disadvantages do however arise in the known methods and apparatus. The apparatus previously described in Fig. 1 may be extended to allow multiple I/O channel bus master scatter gather operations. Clearly, by adding additional physical region descriptor (PRD) tables, "n" dedicated PRD tables may exist, one for each I/O channel to be managed by the I/O control device, regardless of each I/O channel's protocol and function; so long as the I/O channel and its associated I/O devices are capable of DMA transfers. Furthermore, from one to "n" data managers may exist, whose function is to transfer data between the I/O devices residing on the "n" I/O channels and main memory. Prior art implementations for performing these data transfers typically dedicate a single data manager to each I/O channel because the methods used to control the data managers are not successful in maintaining maximum possible I/O bandwidth between the I/O devices and main memory.

Fig. 3 shows single and dual data manager performance scenarios for a dual channel scatter gather I/O controller using the prior apparatus and methods just described. It should be noted that the I/O channel is idle when the active I/O device's buffer has been emptied by the I/O channel DMA. As will be appreciated, the I/O channels shown in Fig. 3 perform at a bandwidth of twice the media rate of the I/O device.

Historically, I/O devices, due to the nature of their technology and their mechanical or network delays, are able to sustain data transfers at only 1/16 to 1/2 of the I/O bus bandwidth. Most I/O devices contain buffer memories (32 to 256 K bytes in size) to compensate for their low media or network bandwidths. Thus, once the

slow media rate has partially filled the device's buffer memory, I/O bus transfers may progress at data rates dictated by the faster bus bandwidth rather than rates dictated by the slower media or network performance. However, once the buffer space has been exhausted due to a transfer which is larger than the buffer size, transfer bandwidth reverts back to the media or network bandwidth. It should also be noted in Fig 3 that the prior method of performing DMAs does not efficiently utilize I/O device buffer memory during scatter gather operations. In fact, substantial media idle time and I/O bus idle time is incurred when using prior art methods with a single data manager and two I/O channels. The single data manager mechanism compromises the aggregate I/O performance by requiring the completion of one I/O channel's PRD table prior to starting a bus master operation on the alternate channel. Dual data manager I/O controllers overcome these problems, but are expensive to implement due to their large FIFOs and complex DMA engines.

In spite of numerous recent I/O related developments and proposals in the PC industry, additional capabilities are needed for achieving maximum I/O performance in multi-tasking, multi-threaded operating systems, at reasonable cost.

The present invention seeks to provide a data processing method and apparatus having advantage over known apparatus and method.

According to one aspect of the present invention there is provided a method of transferring blocks of data from a plurality of input/output devices in a data processing system, characterised by the step of

scheduling DMA block transfers based on the availability of data from the plurality of input/output devices.

According to another aspect of the present invention there is provided data processing apparatus comprising central processing means, memory means, input/output controller means coupled between said central processing means and said memory means, and characterised by a plurality of input/output devices coupled to said input/output controller means, said input/output devices each having a buffer, wherein said input/output controller comprises means for transferring data from said input/output device to said memory means based upon availability of data in at least one of said buffers.

According to a further aspect of the present invention there is provided a method of performing DMA data transfers in a data processing system, characterised by the steps of interleaving blocks of data based upon locations of scatter gather physical region descriptor entries maintained in a table.

The present invention therefore advantageously provides an apparatus and method for improving the input/output performance of a computer system under the control of a multi-tasking, multi-threaded operating system. In particular, the invention provides an apparatus

and method to chain contiguous DMA scatter gather sub blocks of a PRD table for channel 0 with contiguous DMA scatter gather sub blocks of a PRD table for channel 1, using a single data manager, while maintaining maximum media bandwidth. DMA block transfers are scheduled based on the availability of data from the I/O device's buffer memory, thus minimizing both media and network idle time as well as minimizing I/O buses and their associated devices is obtained. The apparatus and method thus provides significant performance advantages over prior techniques having two I/O channel systems implemented with a single data manager.

The present invention can therefore provide an apparatus and method for further improving computer systems I/O performance at minimal cost.

The apparatus and method provides a means for preempting outstanding and in-progress DMA requests to an I/O device on a given I/O channel when a higher priority request is encountered for a device on the same channel or for the same device.

The present invention is advantageous in providing an improved data processing system.

Also, the present invention can provide an improved I/O operation in a data system, and provide improved DMA operations in a data processing system.

Further, the present invention can advantageously also provide improved DMA performance, at minimal incremental cost, in a data processing system.

In a homogeneous programming interface environment as described previously, the invention can be readily extended to control "n+i" I/O channels using "n" data managers.

Further, the apparatus can provide a performance efficient method for software drivers to make use of preemptive scheduling techniques found in newer multi-tasking, multi-threaded operating systems. Preemptive support lessens the response time of performance critical I/O requests.

The invention is described further hereinafter, by way of example only, with reference to the accompanying drawings in which:

Fig. 4 is a block diagram of an apparatus for performing single data manager controlled bus master scatter gather DMA operations in a two I/O channel system;

Fig. 5 is a diagram of the method for performing single data manager controlled bus master scatter gather DMA operations in a two I/O channel systems;

Fig. 6 illustrates a method of merging concurrent interrupts;

Fig. 7 illustrates a method of arbitration and selection by a data manager.

Fig. 8 contains performance analysis diagrams for single data manager, dual I/O channel operations;

Fig. 9 contains diagrams showing possible modes of operation supported by the current invention;

Fig. 10 is a block diagram of particular operating environment of the present invention, using a PCI/Dual IDE I/O channel control device; and

Fig. 11 is a block diagram of an alternate operating environment of the present invention, using a PCT/Dual IDE/ethernet I/O channel control device.

Fig. 4 illustrates a particular apparatus embodying the present invention. An input/output control device 120 contains a single data manager 124 and I/O interface logic 126 which interfaces to two IDE I/O channels 127 and 129. A main memory 130 contains two physical region descriptor tables 137 and 139 for controlling bus master DMA transfers on bus/channel 0 and bus/channel 1. A single PRD table is dedicated to each I/O channel, as shown at 121 and 123. The data manager 124 contains the necessary logic to chain DMA scatter gather blocks specified by PRD table entries between channels and between PRD tables in an alternating fashion. The data manager also contains the necessary logic to determine whether or not a target device on either channel contains DMA data ready for transfer. PRD tables 137 and 139 are not "swapped" by the data manager unless the alternate channel is prepared to transfer data.

Fig. 5 illustrates a particular method embodying the present invention and used for performing I/O operations. At steps 142 and 143, the CPU prepares a physical region descriptor (PRD) table in main memory. Each PRD table entry is 8 bytes long and consists of an address pointer to the starting address and the transfer count of the memory buffer to be transferred. The PRD table may contain up to 8192 PRD table entries, with each entry specifying a single, contiguous block of data to be transferred between the I/O channel device and main memory. Sequential PRD table entries may specify that a block of data be transferred to or from any available location in memory. The two PRD tables may be prepared concurrently or sequentially.

At steps 144 and 145, the CPU writes the starting address of the main memory physical region descriptor table to the I/O control device's descriptor table pointer register. The CPU writes the I/O controller's DMA transfer direction bit and clears appropriate I/O controller interrupt bits and error flags resulting from previous transfers. The CPU may write either descriptor table pointer register at any time so long as the table to be written is not an active bus master.

At step 146 or 147 (depending upon which channel is being used), the CPU issues the appropriate DMA command to an I/O device. A second I/O device on the

alternate channel may be issued an additional DMA command concurrently with the previously mentioned I/O device's DMA operation.

At step 148 or 149 (depending upon which channel is being used), the CPU writes a start scatter gather command to one channel of the I/O controller. A second I/O control devices may be issued the start scatter gather command concurrently with the alternate channel's bus master operation.

As shown in the remaining steps of Fig. 5, the I/O controller performs a bus master DMA operation on the first contiguous DMA block transfer available from either active I/O device, as specified by the associated PRD table entry for that channel. Following the above contiguous DMA block transfer, the data manager transfers a block from the alternate channel if data is available from the I/O device on that channel. If data is not available from the alternate channel, and data is available from the original channel, the data manager transfers a block from the original channel. Transfers continue in this manner so long as data is available on one of the two channels, and the current PRD tables have not been exhausted. If neither channel has data available for transfer, the data manager waits for the first available data transfer opportunity. Once data is available from either channel, the alternating channel DMA flow described above is restarted.

At steps 156 or 157 (depending upon which channel is being used), the I/O device determines the transfer status, then signals that the requested data transfer is complete. The I/O controller signals the processor when all of the requested data has traveled through the I/O controller's internal FIFOs and has been written successfully to either main memory or the I/O device. An I/O device on the alternate channel will continue operation following completion of a PRD table by the original channel.

This particular method provides the capability of non-destructively preempting either bus master at any time during the transfer process. Two types of preemption exist, a preemption during mechanical or network delays prior to data transfer across the I/O bus, and preemption during active DMA data transfers. The first preemption (i.e. prior to data transfer) is detected at the 'Yes' exit point from block 158 of Figure 5. The second preemption (i.e. during active DMA data transfers) is detected at the 'Yes' exit point from blocks 154 and 155 of Fig. 5.

During preemption, the process/program clears the start scatter gather command from the I/O control device for the preempted channel and evaluates the appropriate descriptor table pointer.

If the pointer contains the address of the last PRD table entry from a previously executed PRD table, then the preemption is of the "prior to DMA" type. The process/program clears the active I/O device command, sends a new DMA command to the desired I/O device, rewrites the PRD tables, saves the current descriptor

pointer register, rewrites the descriptor pointer register, and sends a new start scatter gather command to the I/O control device. When the preempting PRD table has completed its transfer, the previously saved descriptor pointer register is reloaded following the reissuing of the DMA command to the preempted I/O device.

If the pointer contains the address of any PRD table entry from the current PRD table, then the preemption is of the "during DMA" type. The process/program writes a "stop after current table entry" command to the I/O control device. The data manager will complete the current PRD table entry's DMA, then interrupt the processor.

The process/program clears the active I/O device command, sends a new DMA command to the desired I/O device, rewrites the PRD tables, saves the current descriptor pointer register, rewrites the descriptor pointer register, and sends a new start scatter gather command to the I/O control device. When the preempting PRD table has completed its transfer, the previously saved descriptor pointer register is reloaded following the reissuing of the abbreviated DMA command to the preempted I/O device.

Merging of concurrent interrupts can be achieved and as previously described, multiple contiguous DMA blocks, controlled by PRD table entries in a scatter gather PRD table, are transferred to or from contiguous memory space of a specific I/O device. Following the final data transfer controlled by the last entry in a PRD table, the processor is interrupted either by the I/O device or the I/O controller depending on the I/O controller's current configuration. "n" concurrent data transfer operations for multiple I/O devices on multiple I/O channels complete their transfers based on time scales that are not deterministic in nature. This characteristic of concurrent I/O channel operation requires a mechanism to associate a given interrupt with the responsible I/O device, channel and operation. Furthermore, concurrent I/O operations could generate multiple, overlapping interrupts requiring either a hardware interrupt queuing logic, or complex interrupt handling software for the host processor.

As illustrated in Fig. 6, the present invention can control interrupt generation through selective arbitration. When a scatter gather transfer reaches its final PRD table entry, a processor interrupt is eminent, and will occur following the successful completion of data transfer controlled by that table entry. Other channels' concurrent scatter gather DMA transfers are allowed to proceed through their PRD table so long as they are not executing a final block transfer PRD entry themselves. If a second final PRD table entry transfer is encountered, its operation is deferred until the first channel's transfer completion interrupt has been serviced by the host processor. If more than two active I/O channels exist, and both previous channels have reached last table entry status, the third I/O channel will proceed with scatter gather DMA operations so long as that channel doesn't

encounter an eminent end of table condition prior to the CPU servicing anticipated or outstanding previous data transfer completion interrupts.

Thus, the described techniques provide discrete, non-overlapping interrupts to signal the completion of each I/O block transfer request to the host processor. This capability minimizes the complexity and improves the performance of the host processor's interrupt handling software, when managing concurrent I/O transfers. Prior methods for handling concurrent interrupts required interrupt queues or complex interrupt handling software running on the host CPU.

First order data manager selection and arbitration is based on the capability of a given I/O device to transfer data at any given point in time. Typical I/O devices contain internal buffers which allow data to be transferred to or from the I/O device at maximum I/O channel bandwidth rather than at the slower media or network bandwidths. Fig. 7 illustrates how data transfer is selected/enabled. During I/O read operations, the I/O device requests data transfers to the I/O channel when internal memory buffers have been filled to a preconfigured "high watermark" level, ensuring that the resulting DMA to the I/O channel efficiently utilizes the available I/O bus bandwidth. During I/O write operations, the I/O device requests data transfers from the I/O channel when internal memory buffers are at a "low watermark" level; once again allowing efficient utilization of the I/O bus bandwidth. First order data manager channel selection is based directly on an I/O device's ability to transfer a block of data at any given point in time. This method is efficient due to the selection process being controlled by actual data availability from the I/O device during I/O reads, or buffer availability within the targeted I/O device during I/O writes. This method is much more cost efficient to implement because the anticipatory logic, bus monitoring logic and complex arbitration logic or software or alternative approaches is not required.

Thus, in summary, first order DMA scheduling is on an "as needed" basis depending upon whether or not the targeted I/O devices have transferred enough data from (or to) media or the network to partially fill (or empty) the I/O device buffer. The apparatus allows various second order arbitration and scheduling techniques to be used for scheduling bus master DMA operations among "n" data managers for "n+i" I/O channels when multiple I/O channels meet the first order constraints for bus mastership, while seeking the use of a data manager which has recently become idle. Well known arbitration algorithms such as round robin, first come first serve, least recently used, most recently used and random can be used for second order arbitration.

Fig. 8 illustrates that the preferred method described hereinabove efficiently utilizes the buffering capabilities of the I/O devices as well as the bandwidth of the two I/O channels. It should be noted that, similar to the assumptions given with respect to Fig. 3, the I/O channel is idle when the active I/O device's buffer has

been emptied by the I/O channel DMA. The I/O channels illustrated Fig. 6 perform at a bandwidth of twice the media rate of the I/O device.

A comparison with Fig. 3 shows that the single data manager solution described here performs almost as well as the previously described dual data manager solution, and significantly out performs the previously described single data manager, dual I/O solution. Furthermore, the process/program can tune the I/O transfer performance by modifying the transfer size and number of the PRD table entries based on the I/O devices' buffer size, media rate and priority.

Typical operating modes allowed by the present invention are illustrated in Fig. 9. A single channel operation is shown at 200, where Channel 0 transfers blocks A, B, C and D, followed by Channel 1 transfer of Blocks E, F, G and H.

Concurrent channel operation is shown at 202, where the transfer of blocks A, B, C and D from Channel 0 are interleaved with the transfer of blocks E, F, G and H from Channel 1. Although individual blocks are not transferred concurrently, but rather interleaved, the overall data transfer can be viewed as being concurrent.

Concurrent channel operation with Channel 0 preemption is shown at 204, where data from another device on a particular channel (in this example blocks I and J from a device 1 on Channel 0) preempts the transfer of data from a device on the same channel (in this example, blocks A, B, C and D from device 0 on Channel 0). The blocks are initially transferred in a manner similar to that illustrated with the normal concurrent channel operation of 202. However, prior to transferring Block C, a preemption on channel 0 is received. The process previously described is then performed, including rewriting the descriptor pointer register. Channel 1 then continues DMA transfer of its next block of data (block G). Channel 0 then transfers a block of data (block I) from the preemptive second device on Channel 0. Channel 1 then continues DMA transfer of its next block of data (block H). Channel 0 then transfers the next block of data (block J) from the preemptive second device on Channel 0. As this second device has no more data to transfer, the original PRD for channel 0 is restored, and data transfer from the first device on channel 0 continues with the transfer of blocks C and D.

Concurrent channel operation with Channel 1 preemption is shown at 206. This operates in an analogous fashion to that just described with respect to the preemptive Channel 0 operation of 204 and so further discussion is unnecessary.

Fig. 10 illustrates the I/O Control Device 120 used in the preferred system environment. As shown, the preferred environment is for use in interfacing two channels of an IDE bus with a PCI bus. IDE buses and PCI buses are well known in the art. The I/O control device 120 allows for a more efficient DMA transfer in such an environment, as previously described.

Fig. 11 illustrates an alternate system environment

to that of Fig. 10. The techniques described herein can be extended to support two channels of an IDE bus in addition to an ethernet channel. In order to achieve such expansion, an additional descriptor pointer register 125 for channel 2 is added. An additional DMA engine 225 is also added to provide specific DMA interface support for the additional channel. This DMA engine would operate in similar fashion to that illustrated in Figure 5, but only having flow control for managing a single channel (as opposed to the two channels shown in Fig. 5).

The above described apparatus and method provides a mechanism for performance tuning DMA transfers between main memory and I/O devices, based on the media or network bandwidth and buffer size of the I/O device. Slower I/O devices (i.e. CDROMS, tape drives or ethernet network devices) are controlled using a larger number of scatter gather PRD table entries having a smaller block transfer size. Faster I/O devices (i.e. hard disk drives or fiber channel switch fabric ports) are controlled using a smaller number of scatter gather PRD table entries having a larger block transfer size. Likewise, I/O devices having smaller buffer sizes are controlled using a larger number of scatter gather PRD table entries having a smaller block transfer size. I/O devices having larger buffer sizes may be controlled using a smaller number of scatter gather PRD table entries having a larger block transfer size. Software may easily tune the I/O controller to provide balanced, high aggregate bandwidth system performance.

It should be appreciated that the present invention is not restrictive to the details of the forgoing embodiments.

Claims

1. A method of transferring blocks of data from a plurality of input/output devices in a data processing system, characterised by scheduling DMA block transfers based on the availability of data from the plurality of input/output devices.
2. Data processing apparatus comprising
 - central processing means (34),
 - memory means (130),
 - input/output controller means (120) coupled between said central processing means (34), and characterised by a plurality of input/output devices (127, 129) coupled to said input/output controller means (120), said input/output devices (127, 129) each having a buffer, wherein said input/output controller comprises means for transferring data from said input/output device (127, 129) to said memory means (120) based upon availability of data in at least one of said buffers.

3. A method of performing DMA data transfers in a data processing system, characterised by the step of interleaving blocks of data based upon locations of scatter gather physical region descriptor entries maintained in a table.

FIG. 1

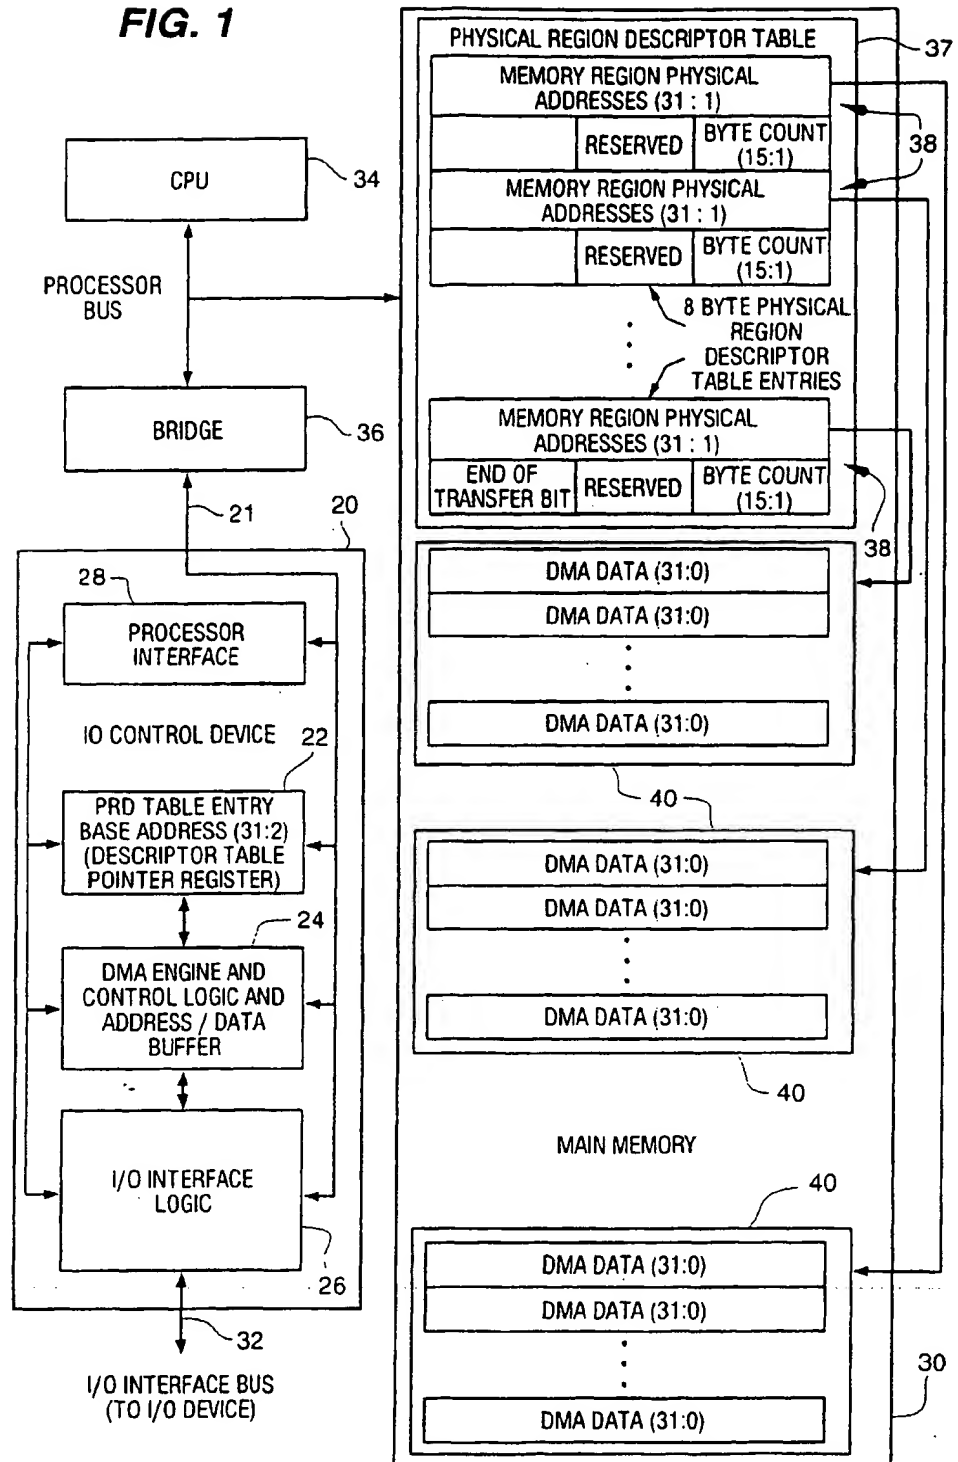


FIG. 2



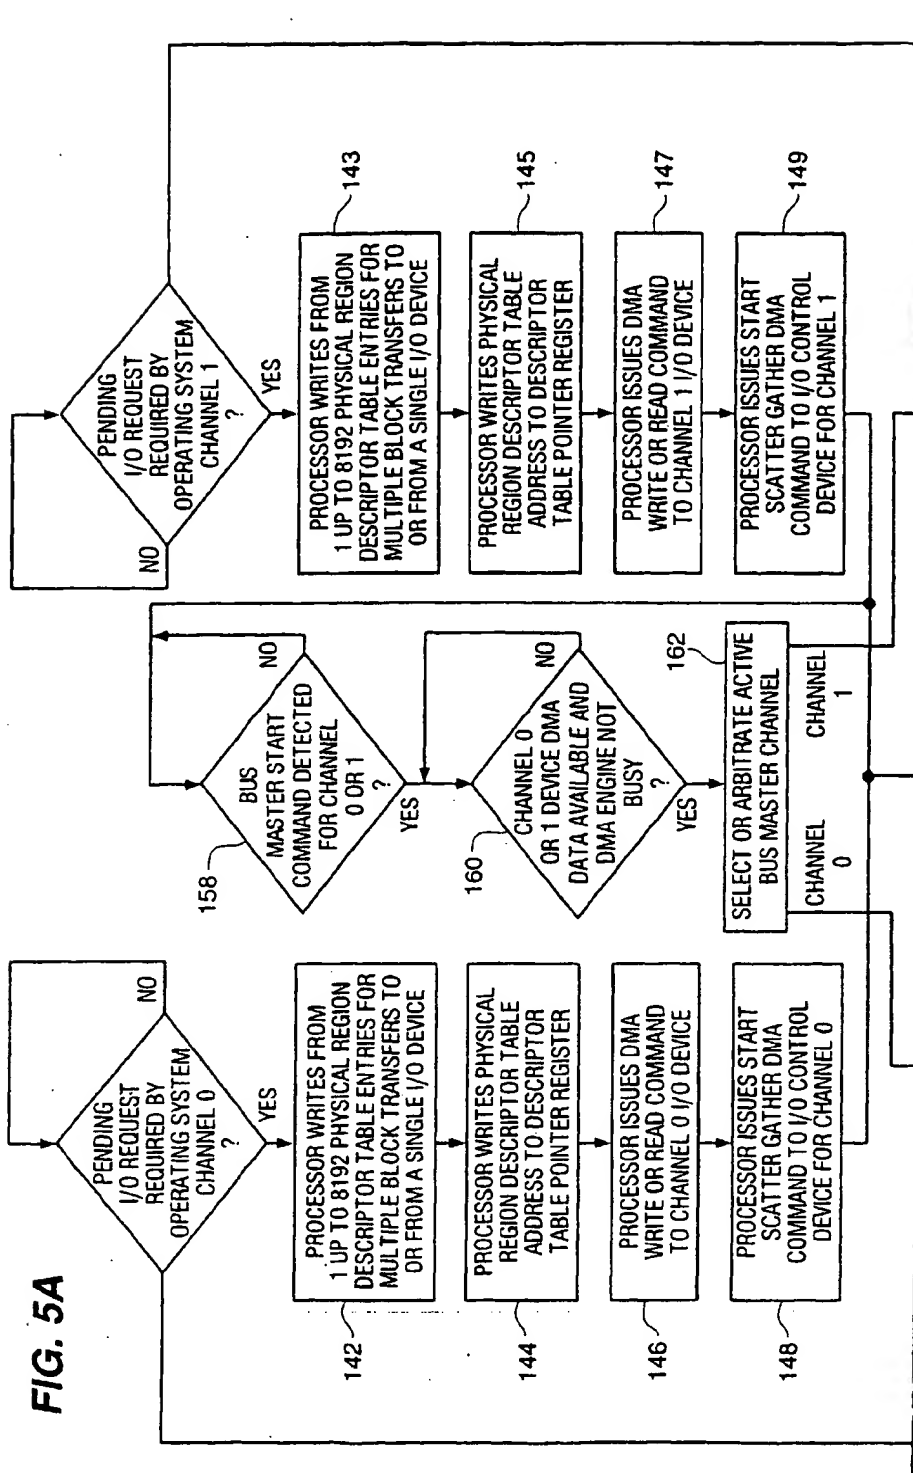




FIG. 3

FIG. 4





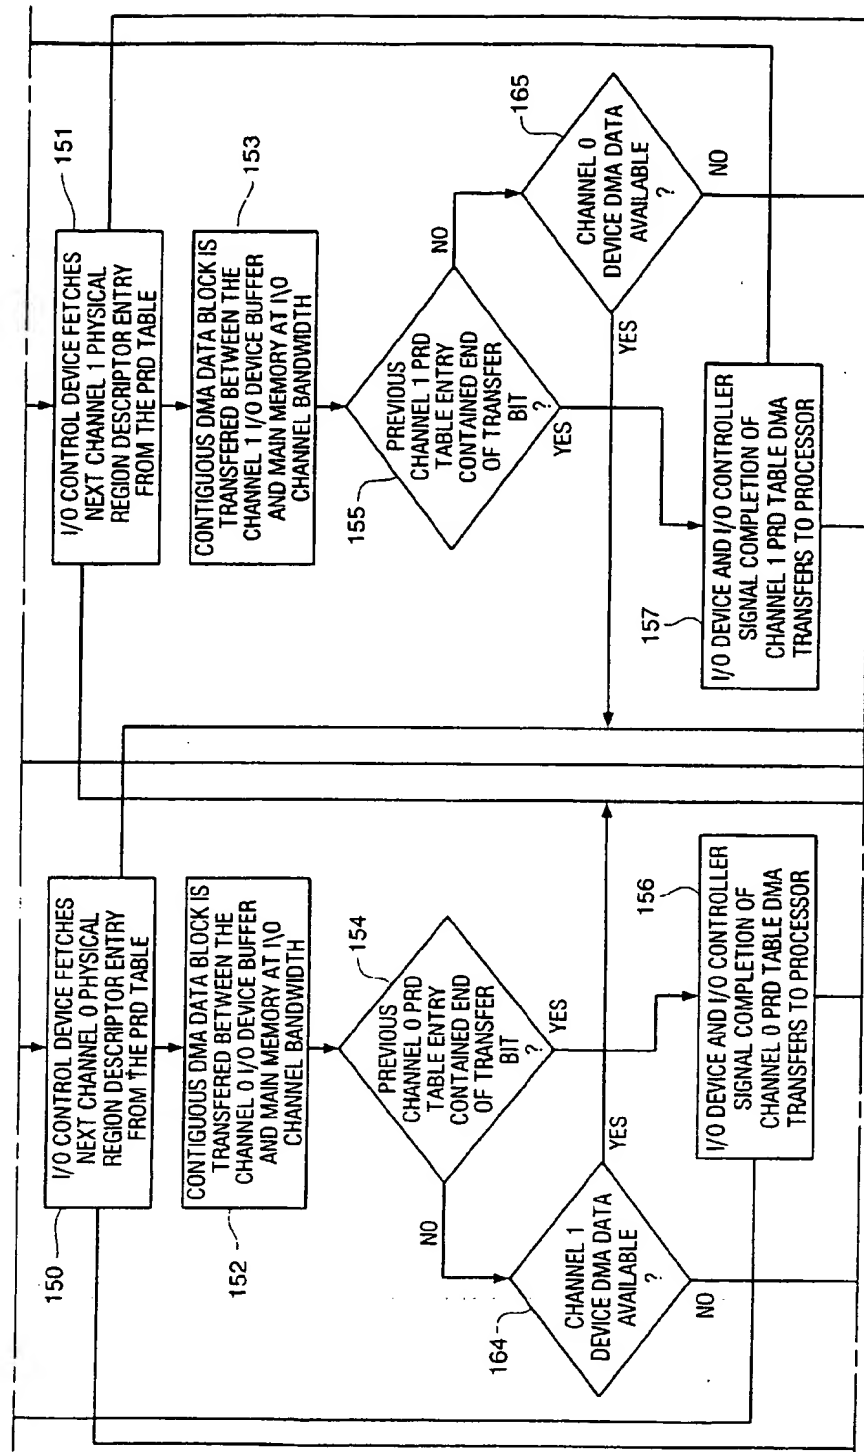


FIG. 5B




FIG. 5C

FIG. 6A

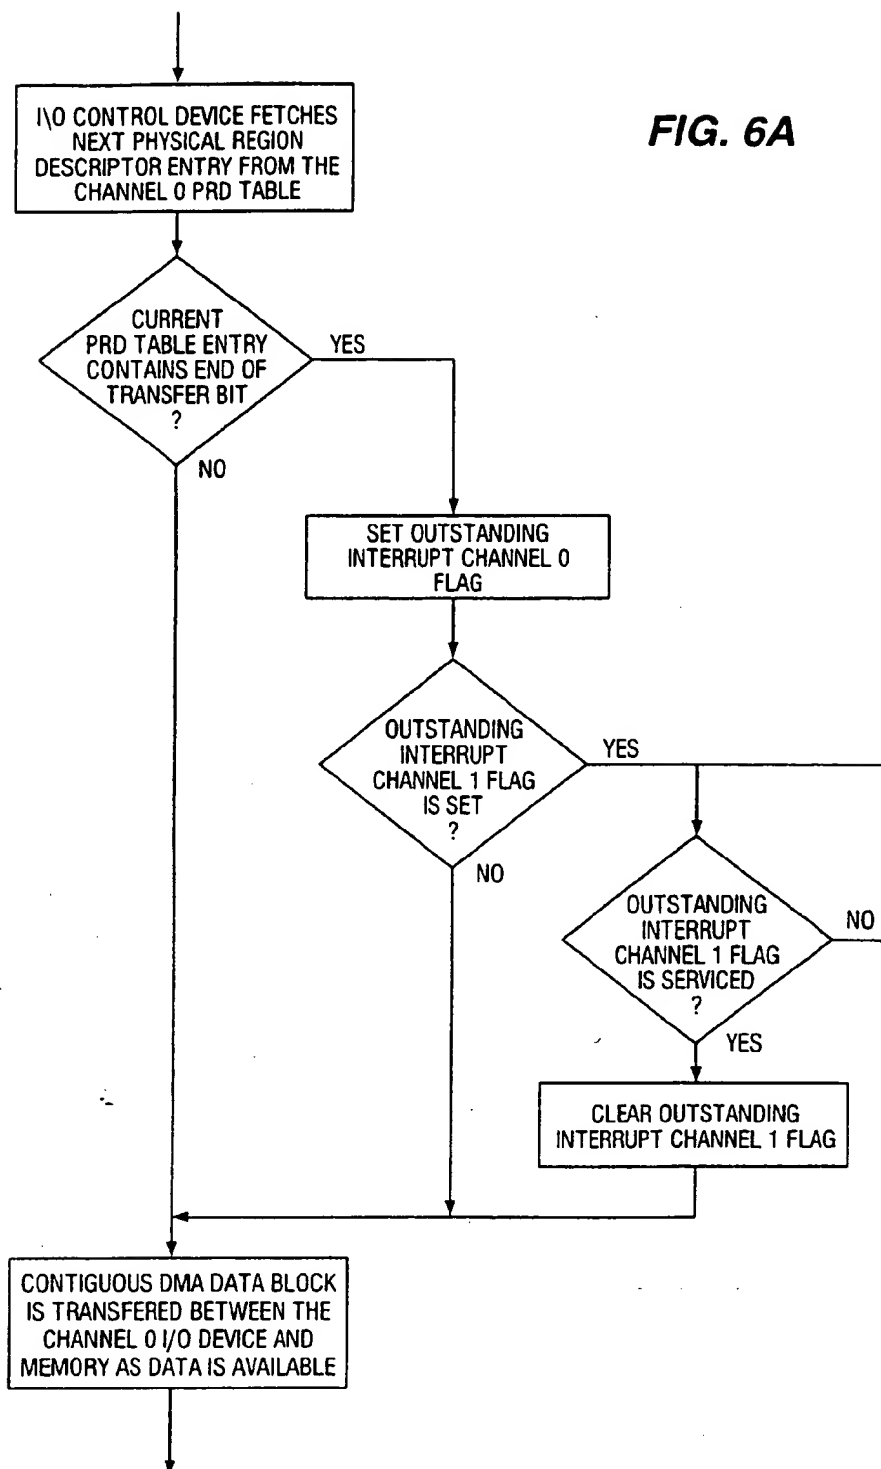


FIG. 6B

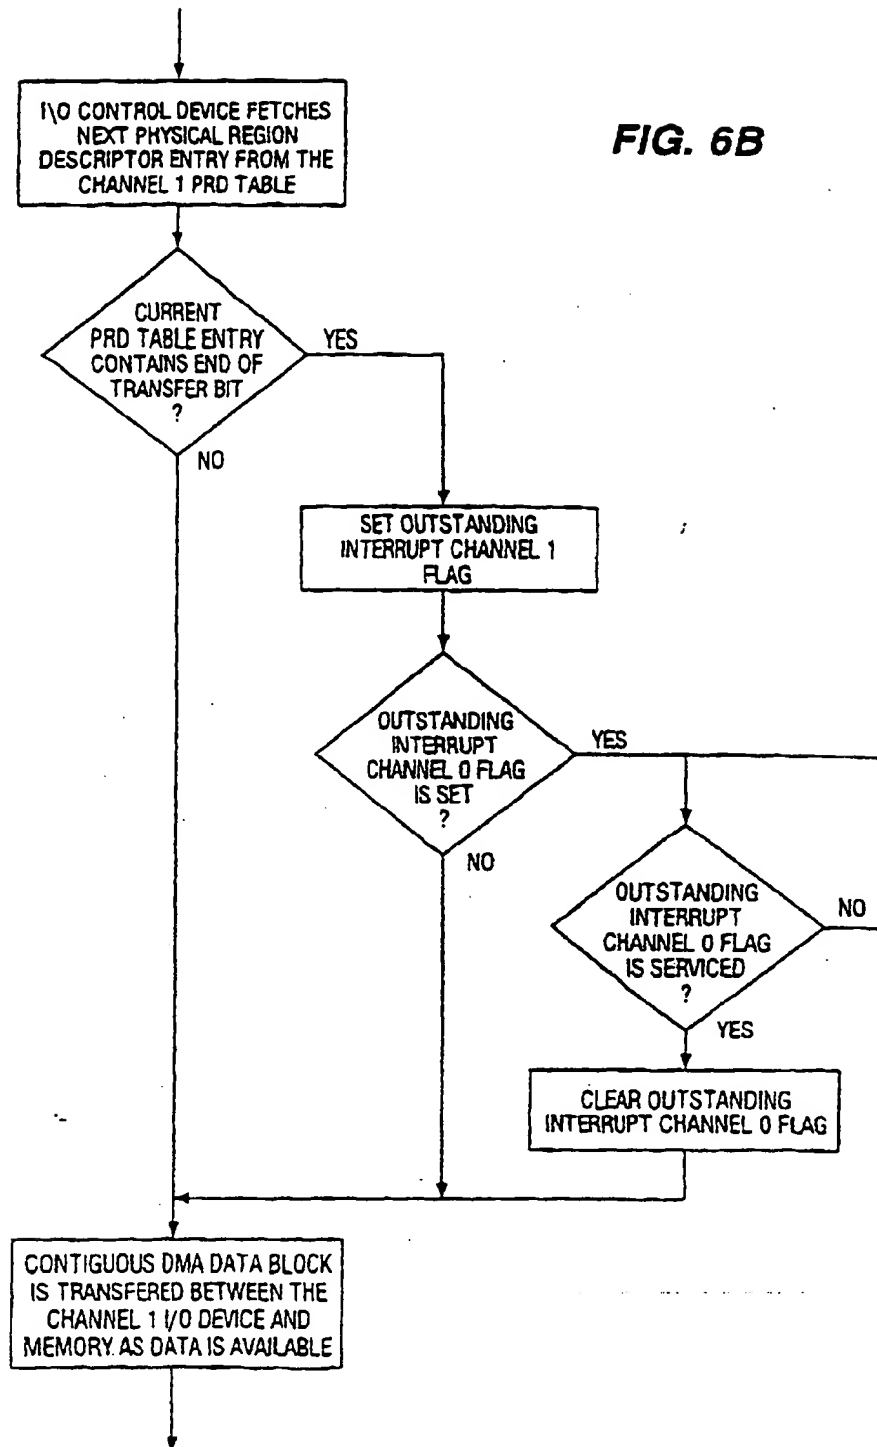
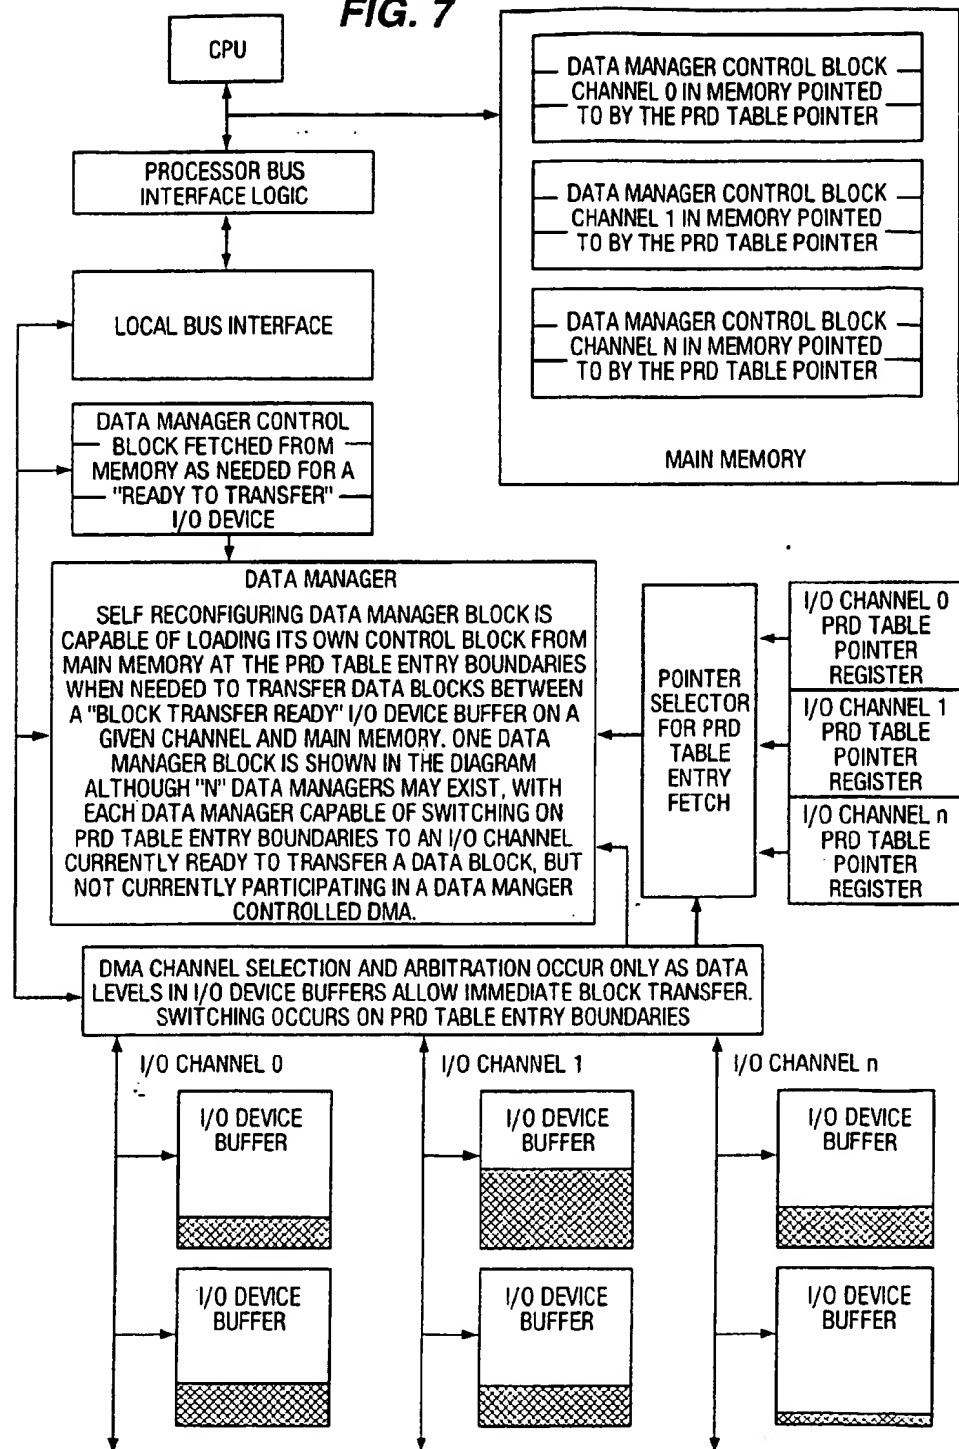




FIG. 7






FIG. 8

FIG. 9




FIG. 10





FIG. 11





European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 96 30 1799

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
X	EP-A-0 537 401 (INTERNATIONAL BUSINESS MACHINES CORPORATION) * column 9, line 19 - column 11, line 2 * * abstract; figure 2 * ---	1,2	G06F13/28
X	IBM TECHNICAL DISCLOSURE BULLETIN, vol. 37, no. 1, January 1994, NEW YORK US, pages 337-342, XP002004997 "DMA controller channel interlocking" * the whole document * ---	1-3	
A	US-A-4 831 523 (LEWIS ET AL.) * column 5, line 4 - line 40; figure 2 * ---	1,2	
A	WO-A-93 06553 (INTERNATIONAL BUSINESS MACHINES CORPORATION) * abstract * ---	1,2	
A	EP-A-0 317 481 (INTERNATIONAL BUSINESS MACHINES, INC.) * page 6, line 45 - page 7, line 15 * ---	1,2	
A	US-A-4 805 137 (GRANT ET AL.) * column 3, line 45 - line 63; figure 6 * ---	3	
A	EP-A-0 530 543 (STRATUS COMPUTER, INC.) * abstract * ---	3	
A	EP-A-0 549 924 (INTERNATIONAL BUSINESS MACHINES CORPORATION) * column 8, line 45 - column 9, line 30; figure 1 * ---	3	
		-/--	
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 7 June 1996	Examiner Jones, H
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons</p> <p>& : member of the same patent family, corresponding document</p>			

EPF FORM 1501 (04/92) (P04C01)



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 96 30 1799

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
A	IBM TECHNICAL DISCLOSURE BULLETIN, vol. 38, no. 2, February 1995, NEW YORK US, pages 477-480, XP002004998 "priority scheme for arithmetic logic unit and dataflow usage by P1394 isochronous hardware" * the whole document * -----	1-3	
			TECHNICAL FIELDS SEARCHED (Int.Cl.6)
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 7 June 1996	Examiner Jones, H
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document		T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document	

EPO FORM 1503 (01/92) (P0401)